

# High Performance Computing at Stellenbosch University

Gerhard Venter



# Outline

- 1 Background
- 2 Clusters
- 3 SU History
- 4 SU Cluster
- 5 Using the Cluster
- 6 Examples



# What is High Performance Computing?

## Wikipedia

High-performance computing (HPC) uses supercomputers and computer clusters to solve advanced computing problems. The term is most commonly associated with computing used for scientific research.

- Supercomputers
  - Bleeding edge technology
  - Typically one-of-a-kind, custom design
  - EXPENSIVE
- Computer cluster
  - Group of closely linked computers working together in parallel
  - Number of ordinary (off the shelf) processors (CPU's)
  - Relatively inexpensive



# What is High Performance Computing?

## Wikipedia

High-performance computing (HPC) uses supercomputers and computer clusters to solve advanced computing problems. The term is most commonly associated with computing used for scientific research.

- Supercomputers
  - Bleeding edge technology
  - Typically one-of-a-kind, custom design
  - **EXPENSIVE**
- Computer cluster
  - Group of closely linked computers working together in parallel
  - Number of ordinary (off the shelf) processors (CPU's)
  - Relatively inexpensive



# What is High Performance Computing?

## Wikipedia

High-performance computing (HPC) uses supercomputers and computer clusters to solve advanced computing problems. The term is most commonly associated with computing used for scientific research.

- Supercomputers
  - Bleeding edge technology
  - Typically one-of-a-kind, custom design
  - **EXPENSIVE**
- Computer cluster
  - Group of closely linked computers working together in parallel
  - Number of ordinary (off the shelf) processors (CPU's)
  - Relatively inexpensive



# What is a Computer Cluster?

## Wikipedia

A computer cluster is a group of linked computers, working together closely so that in many respects they form a single computer. The components of a cluster are commonly, but not always, connected to each other through fast local area networks. Clusters are usually deployed to improve performance and/or availability over that provided by a single computer, while typically being much more cost-effective than single computers of comparable speed or availability.


- Linux is the OS of choice



# What Does a Linux Cluster Look Like?



# What Does a Linux Cluster Look Like?






# What Does a Linux Cluster Look Like?



# What Does a Linux Cluster Look Like?



# Typical Linux Cluster Layout



# History at Stellenbosch University

- M&M had some money and the need for a fast computer
- Number of small “white box” clusters exist at SU
  - Computer Science
  - Physics
  - Biochemistry
  - Chemistry and Polymer Science
- M&M decided campaign for a central facility that is accessible to all researchers
- Final funding from both M&M and the University



# History at Stellenbosch University ...

- Initial thought was roughly 70 “white box” PC’s in a room
- System housed at M&M
- With the help of IT we decided to go with:
  - Rack mounted system
  - System housed in IT’s data centre
  - IT provides system admin
  - IT manages (and upgrades) facility free of charge for all researchers



# Stellenbosch University Linux Cluster

- Worked with IT to start an official tender process
- Quickly realised that it is extremely difficult to provide specifications for the system
  - General lack of knowledge
  - Fast changing environment
  - No single solution that will satisfy everybody in a general environment
- The tender process when through two cycles and the final bid went to Sun Microsystems
- Job scheduler is Sun Grid Engine (SGE)



# Stellenbosch University Linux Cluster ...

rhasatsha(the *rha* is pronounced as *gaan* in Afrikaans)

a clever person/object; highly intelligent; something that acts promptly; a wide awake person/object who/that is always on the spot; a versatile person/object that can tackle anything successfully.

- Currently the system is still referred to as  
`head001.sun.ac.za`



# Stellenbosch University Linux Cluster ...

rhasatsha(the *rha* is pronounced as *gaan* in Afrikaans)

a clever person/object; highly intelligent; something that acts promptly; a wide awake person/object who/that is always on the spot; a versatile person/object that can tackle anything successfully.

- Currently the system is still referred to as  
`head001.sun.ac.za`





# Rhasatsha: Specifications

## Rhasatsha Cluster



CPU: 168 cores @ 2.83Ghz  
 Memory: 336GB  
 I/O: 12.1TB  
 Power Idle: 5.6 kW  
 Power Peak: 7.5 kW  
 Theoretical performance: 2 Tflop

Linpack score: coming soon

### Job Manager Sun Grid Engine 6.0



### Head Node Specification (1)

Model: X4150  
 Processors: 1 x E5410 (2.33Ghz) Harpertown 45nm Quad-core  
 Memory: 4GB FBDIMM  
 I/O: 2 x 146GB 10K rpm 2.5" SAS

### Compute Node Specification (21)

Model: X4150  
 Processors: 2 x E5440 (2.83Ghz) Harpertown 45nm Quad-core  
 Memory: 16GB FBDIMM  
 I/O: 2 x 146GB 10K rpm 2.5" SAS

### Inter-Connect Model: Extreme Summit 4x0 Ports: 24 x 1Gbit Ethernet



### Storage

Model: 2530  
 Controllers: 1 x 3 port (3Gbit) SAS/SATA  
 Memory: 512MB cache  
 I/O: 12 x 500GB 7200 rpm 3.5" SATA



# Rhasatsha: Photos



# How do I use the Cluster?

- Work through the network
- Need to make a `ssh` connection to the head node
- Copy data using `scp`
- Can connect from either Linux or Windows
  - Most Linux have both `ssh` and `scp` clients installed by default
  - Windows require the addition of a `ssh` and a `scp` client



# A Typical Session

- Connect to the head node
  - Access your own account/home directory
  - No direct access to the compute nodes
- Setup your job
  - Copy/create your input data and source code
  - Port your source code or use a commercial code
  - Specify a submit script for the scheduler
- Submit your job
- Remember that you are working on a Linux system
  - Windows programs do not run on Linux
  - Linux file system is case sensitive



# Some Useful Commands

- A Linux few commands for Windows users
- The following commands can be executed in the `ssh` client
- Most other actions can be performed using the `scp` client

## Linux

- `ls -lrt`  
List files
- `pwd`  
Present work directory
- `cd <dir name>`  
Change directory

## SGE

- `qsub <script name>`  
Submit your job
- `qstat -f`  
Check status of your job
- `qdel <job number>`  
Delete a running job

# Why do I Submit a Job?

- To run a job on the compute nodes - submit the job to SGE
- The job and the required resources are specified in a script
- The scheduler decides when and where to run the job
- The scheduler can handle different types of jobs
- The most common jobs can be classified as:
  - Serial jobs
  - Parallel array jobs
  - OpenMP parallel jobs
  - MPI (OpenMPI) parallel jobs



# Serial Jobs

- Run a serial (single core) instance of a job
- No parallel computing is involved

## Serial Script

```
#!/bin/bash  
#$ -cwd -j y  
/vrand/bin/genesis beam.dat
```



# Array Jobs

- Simply run multiple instances of the same executable
- Each instance will typically process different input files
- No communication between instances
- Useful when doing a parametric or a DOE study

## Array Script

```
#!/bin/bash
#$ -cwd -j y
#$ -t 1-10
/vrand/bin/genesis beam_${SGE_TASK_ID}.dat
```





# OpenMP Jobs

- Useful for some commercial software (e.g. DYTRAN, GENESIS)
- Useful for creating your own code in a shared memory parallel (SMP) environment
- Can automatically create parallel code using compiler flags
- Limited to multiple processors/cores on a single machine

## OpenMP Script

```
#!/bin/bash
#$ -pe openmp 8
#$ -cwd -j y
export PATH=/apps/msc/dytran2007r1/bin:$PATH
dytran bat=no jid=beam ncpus=$NSLOTS
```



# MPI Jobs

- Useful for some commercial software (e.g. Fluent)
- Useful for creating your own code in a shared or distributed memory parallel (DMP) environment
- Can run on a single or multiple machines
- No automatic generate of parallel code

## OpenMP Script

```
#!/bin/bash
#$ -pe openmpi_rr 128
#$ -cwd -j y
export PATH=/apps/Fluent.Inc/bin:$PATH
export LM_LICENSE_FILE=\
/apps/Fluent.Inc/FSLM10.8/license.dat
fluent 3d -g -t$NSLOTS -ssh -i wing.jnl
```



# What if I don't Know Linux?

- You can still access the cluster from Windows
- You need `ssh` and `scp` clients for Windows
  - Putty for the `ssh` client
  - WinSCP for the `scp` client
  - Obtain both from `archive.sun.ac.za`

## Example Problems



# What if I don't Know Linux?

- You can still access the cluster from Windows
- You need `ssh` and `scp` clients for Windows
  - Putty for the `ssh` client
  - WinSCP for the `scp` client
  - Obtain both from `archive.sun.ac.za`

## Example Problems



# Example: Fluent Application

- Commercial CFD software with DMP parallel capabilities
- Very easy to use in the parallel environment
  - Setup your model as normal
  - Create a work directory on the cluster
  - Copy your model to the cluster
  - Create a submit script (work from a sample)
  - Submit your job



# Example: GENESIS Application

- Commercial FEM software with SMP parallel capabilities
- Very easy to use in the parallel environment
  - Setup your model as normal
  - Create a work directory on the cluster
  - Copy your model to the cluster
  - Create a submit script (work from a sample)
  - Submit your job



# Example: MPI Application

- Very simple example that shows message passing
- Master node is identified and sends a greeting to all worker nodes
- Worker nodes receive the greeting, append a response and reply back to the master node
- Master node prints messages from worker nodes as they are received
- Demonstrates the basic nature of a MPI program
  - Write and test the source code with message passing between nodes
  - Compile the source code on the cluster
  - Create a submit script
  - Submit your job



# Conclusions

- HPC computing is a reality at SU
- The Linux cluster is available to all researchers
- Keep in mind that there is no perfect solution for a general purpose cluster
- The SU cluster should provide a real working environment to develop and test new HPC code and applications
- The SU cluster can be used as an easy stepping stone to the CHPC cluster





# How do I get Access?

- The cluster is available to all researchers
- Faculty can request access themselves
- Post-graduate students must apply for access through their study leaders
- To obtain an account, please contact Johann Spies ([jspies@sun.ac.za](mailto:jspies@sun.ac.za)) at IT
  - US number
  - E-mail address
  - Office phone number
  - Study leader for post-graduate students
  - Department
  - Short motivation



# Questions

